

Amendments to the Claims:

Please cancel claims 16 and 21, without prejudice.

Please amend claims 10 and 18.

Please add new claims 23-31.

This listing of claims will replace all prior versions, and listings of claims in the application:

Listing of Claims:

10. **(Currently Amended)** A method comprising:

providing a virtual router (VR)-based switch configured for operation at an Internet point-of-presence (POP) of a service provider, the VR-based switch having a plurality of processing elements;

providing a network operating system (NOS) on each of the plurality of processing elements;

segmenting resources of the VR-based switch between at least a first subscriber of the service provider and a second subscriber of the service provider by:

associating a first plurality of VRs with the first subscriber;

associating a second plurality of VRs with the second subscriber;

mapping the first plurality of VRs onto a first set of one or more of the plurality of processing elements;

mapping the second plurality of VRs onto a second set of one or more of the plurality of processing elements;

wherein a shared processing element of the plurality of processing elements is part of the first set of one or more of the plurality of processing elements and the shared processing element is part of the second set of one or more of the plurality of processing elements;

configuring a first set of customized services including a plurality of firewalling, virtual private networking, encryption, traffic shaping, routing and network address translation (NAT) to be provided by the VR-based switch on behalf of the first subscriber by allocating a first service object group within the first plurality of VRs, the first service object group including a service object corresponding to each service of the first set of customized services and wherein each service object of the first service object group can be dynamically distributed by the NOS to customized processors of the first set of one or more of the plurality of processing elements to achieve desired computational support; and

configuring a second set of customized services including a plurality of firewalling, virtual private networking, encryption, traffic shaping, routing and NAT to be provided by the VR-based switch on behalf of the second subscriber by allocating a second service object group within the second plurality of VRs, the second service object group including a service object corresponding to each service of the second set of customized services and wherein each service object of the second service object group can be dynamically distributed by the NOS to customized processors of the second set of one or more of the plurality of processing elements to achieve desired computational support.

11. **(Previously Presented)** The method of claim 10, further comprising:
 - defining a system VR within the VR-based switch; and
 - the system VR aggregating traffic from the first plurality of VRs and the second plurality of VRs and transferring the aggregated traffic across the Internet.
12. **(Previously Presented)** The method of claim 10, wherein at least one of the first plurality of VRs spans two or more of the first set of one or more of the plurality of processing elements.
13. **(Previously Presented)** The method of claim 11, wherein at least one of the second plurality of VRs spans two or more of the second set of one or more of the plurality of processing elements.

14. **(Previously Presented)** The method of claim 10, further comprising defining a first configured topology among the first plurality of VRs by configuring virtual interfaces (VIs) of the first plurality of VRs to provide desired paths for packet flows associated with the first subscriber and permissible transformations of the packet flows associated with the first subscriber.
15. **(Previously Presented)** The method of claim 14, further comprising defining a second configured topology among the second plurality of VRs by configuring virtual interfaces (VIs) of the second plurality of VRs to provide desired paths for packet flows associated with the second subscriber and permissible transformations of the packet flows associated with the second subscriber.
16. **(Cancelled)**
17. **(Previously Presented)** The method of claim 10, wherein the VR-based switch includes a first server blade and a second server blade and each of the plurality of processing elements are associated with the first server blade or the second server blade, and wherein a VR of the first plurality of VRs terminates links on both the first server blade and the second server blade, and the method further comprises forwarding agents associated with the VR maintaining a replicated forwarding information base.
18. **(Currently Amended)** A method comprising:
 - providing a virtual router (VR)-based switch within a service provider network, the VR-based switch having a plurality of processing elements;
 - providing a network operating system (NOS) on each of the plurality of processing elements;
 - segmenting resources of the VR-based switch between at least a first subscriber of the service provider and a second subscriber of the service provider by:
 - associating a first VR with the first subscriber;
 - associating a second VR with the second subscriber;

mapping the first VRs onto a first set of two or more of the plurality of processing elements;

mapping the second VRs onto a second set of two or more of the plurality of processing elements;

wherein the first set of two or more of the plurality of processing elements and the second set of two or more of the plurality of processing elements have at least one processing element of the plurality of processing elements in common;

configuring a first set of customized services including a plurality of firewalling, virtual private networking, encryption, traffic shaping, routing and network address translation (NAT) to be provided by the VR-based switch on behalf of the first subscriber by allocating a first service object group within the first VR, the first service object group including a service object corresponding to each service of the first set of customized services and wherein each service object of the first service object group can be dynamically distributed by the NOS among processors associated with the first set of two or more of the plurality of processing elements to achieve desired computational support; and

configuring a second set of customized services including a plurality of firewalling, virtual private networking, encryption, traffic shaping, routing and NAT to be provided by the VR-based switch on behalf of the second subscriber by allocating a second service object group within the second VR, the second service object group including a service object corresponding to each service of the second set of customized services and wherein each service object of the second service object group can be dynamically distributed by the NOS among processors associated with the second set of two or more of the plurality of processing elements to achieve desired computational support.

19. **(Previously Presented)** The method of claim 18, further comprising defining desired paths through the VR-based switch for packet flows associated with the first VR by configuring one or more virtual interfaces (VIs) of the first VR.

20. **(Previously Presented)** The method of claim 18, further comprising defining permissible transformations of packet flows associated with the first VR by configuring one or more VIs of the first VR.
21. **(Cancelled)**
22. **(Previously Presented)** The method of claim 18, wherein the VR-based switch includes a first server blade and a second server blade and each of the plurality of processing elements are associated with the first server blade or the second server blade, and wherein the first VR terminates links on both the first server blade and the second server blade, and the method further comprises forwarding agents associated with the first VR maintaining a replicated forwarding information base.
23. **(New)** A method comprising:
 - providing a virtual router (VR)-based switch configured for operation at an Internet point-of-presence (POP) of a service provider, the VR-based switch having a plurality of processing elements and including a first server blade and a second server blade and each of the plurality of processing elements are associated with the first server blade or the second server blade, and wherein a VR of the first plurality of VRs terminates links on both the first server blade and the second server blade;
 - providing a network operating system (NOS) on each of the plurality of processing elements;
 - segmenting resources of the VR-based switch between at least a first subscriber of the service provider and a second subscriber of the service provider by:
 - associating a first plurality of VRs with the first subscriber;
 - associating a second plurality of VRs with the second subscriber;
 - mapping the first plurality of VRs onto a first set of one or more of the plurality of processing elements;

- mapping the second plurality of VRs onto a second set of one or more of the plurality of processing elements;
- configuring a first set of customized services including a plurality of firewalling, virtual private networking, encryption, traffic shaping, routing and network address translation (NAT) to be provided by the VR-based switch on behalf of the first subscriber by allocating a first service object group within the first plurality of VRs, the first service object group including a service object corresponding to each service of the first set of customized services and wherein each service object of the first service object group can be dynamically distributed by the NOS to customized processors of the first set of one or more of the plurality of processing elements to achieve desired computational support;
- configuring a second set of customized services including a plurality of firewalling, virtual private networking, encryption, traffic shaping, routing and NAT to be provided by the VR-based switch on behalf of the second subscriber by allocating a second service object group within the second plurality of VRs, the second service object group including a service object corresponding to each service of the second set of customized services and wherein each service object of the second service object group can be dynamically distributed by the NOS to customized processors of the second set of one or more of the plurality of processing elements to achieve desired computational support; and
- forwarding agents associated with the VR maintaining a replicated forwarding information base.

24. **(New)** The method of claim 23, further comprising:

- defining a system VR within the VR-based switch; and
- the system VR aggregating traffic from the first plurality of VRs and the second plurality of VRs and transferring the aggregated traffic across the Internet.

25. **(New)** The method of claim 23, wherein at least one of the first plurality of VRs spans two or more of the first set of one or more of the plurality of processing elements.

26. (New) The method of claim 25, wherein at least one of the second plurality of VRs spans two or more of the second set of one or more of the plurality of processing elements.
27. (New) The method of claim 23, further comprising defining a first configured topology among the first plurality of VRs by configuring virtual interfaces (VIs) of the first plurality of VRs to provide desired paths for packet flows associated with the first subscriber and permissible transformations of the packet flows associated with the first subscriber.
28. (New) The method of claim 27, further comprising defining a second configured topology among the second plurality of VRs by configuring virtual interfaces (VIs) of the second plurality of VRs to provide desired paths for packet flows associated with the second subscriber and permissible transformations of the packet flows associated with the second subscriber.
29. (New) A method comprising:
 - providing a virtual router (VR)-based switch within a service provider network, the VR-based switch having a plurality of processing elements and including a first server blade and a second server blade and each of the plurality of processing elements are associated with the first server blade or the second server blade, and wherein the first VR terminates links on both the first server blade and the second server blade;
 - providing a network operating system (NOS) on each of the plurality of processing elements;
 - segmenting resources of the VR-based switch between at least a first subscriber of the service provider and a second subscriber of the service provider by:
 - associating a first VR with the first subscriber;
 - associating a second VR with the second subscriber;
 - mapping the first VRs onto a first set of two or more of the plurality of processing elements;

mapping the second VRs onto a second set of two or more of the plurality of processing elements;

configuring a first set of customized services including a plurality of firewalling, virtual private networking, encryption, traffic shaping, routing and network address translation (NAT) to be provided by the VR-based switch on behalf of the first subscriber by allocating a first service object group within the first VR, the first service object group including a service object corresponding to each service of the first set of customized services and wherein each service object of the first service object group can be dynamically distributed by the NOS among processors associated with the first set of two or more of the plurality of processing elements to achieve desired computational support;

configuring a second set of customized services including a plurality of firewalling, virtual private networking, encryption, traffic shaping, routing and NAT to be provided by the VR-based switch on behalf of the second subscriber by allocating a second service object group within the second VR, the second service object group including a service object corresponding to each service of the second set of customized services and wherein each service object of the second service object group can be dynamically distributed by the NOS among processors associated with the second set of two or more of the plurality of processing elements to achieve desired computational support; and

forwarding agents associated with the first VR maintaining a replicated forwarding information base.

30. **(New)** The method of claim 29, further comprising defining desired paths through the VR-based switch for packet flows associated with the first VR by configuring one or more virtual interfaces (VIs) of the first VR.

31. **(New)** The method of claim 29, further comprising defining permissible transformations of packet flows associated with the first VR by configuring one or more virtual interfaces (VIs) of the first VR.